Knowledge Rich Natural Language Queries over Structured Biological Databases

Hasan M. Jamil
Department of Computer Science
University of Idaho
jamil@uidaho.edu

arXiv:1703.10692v1 [cs.DB] 30 Mar 2017

ABSTRACT

Increasingly, keyword, natural language and NoSQL queries are being used for information retrieval from traditional as well as non-traditional databases such as web, document, image, GIS, legal, and health databases. While their popularity are undeniable for obvious reasons, their engineering is far from simple. In most part, semantics and intent preserving mapping of a well understood natural language query expressed over a structured database schema to a structured query language is still a difficult task, and research to tame the complexity is intense. In this paper, we propose a multi-level knowledge-based middleware to facilitate such mappings that separate the conceptual level from the physical level. We augment these multi-level abstractions with a concept reasoner and a query strategy engine to dynamically link arbitrary natural language querying to well defined structured queries. We demonstrate the feasibility of our approach by presenting a Datalog based prototype system, called *BioSmart*, that can compute responses to arbitrary natural language queries over arbitrary databases once a syntactic classification of the natural language query is made.

CCS Concepts

•**Information systems** → **Query optimization; Query languages for non-relational engines;** *Query planning;* Semi-structured data; Web interfaces; •**Human-centered computing** → **Graphical user interfaces;** Web-based interaction; Graph drawings; •**Applied computing** → **Bioinformatics;** *Genomics;* •**Computing methodologies** → *Logic programming and answer set programming;*

1. INTRODUCTION

An overwhelming majority of scientific databases use traditional database query interfaces such as SQL and XQuery, or predesigned graphical query interfaces to grant access to their contents. As the information contents of these databases grow more complex in representation, interpretation and analysis, query interfaces needed to access them are also becoming increasingly complicated. Often time, the only convenient method is a predesigned graphical interface through which all accesses are facilitated even though it severely limits the usefulness of these rich information repositories. Although widely used in Life Sciences, these interfaces do not allow ad hoc or spontaneous queries, or investigative queries in ways a natural language interface (NLI) would allow. An encouraging sign is that NLIs are increasingly being used to provide non traditional query responses in various types of databases such as knowledge repositories [8, 18, 15], text based information repositories [51, 19, 23], biological and clinical databases [62, 24, 52], GIS databases [46] and of course traditional relational databases [65, 47, 48] with limited scopes.

From the standpoint of end users in Life Sciences, flat view of data is perhaps the most acceptable of all formats while some forms of shallow nesting were also welcomed. Therefore, it can be argued that the relational data model fits well with the practice and end user psychology well in this domain. While XML is creeping its way up in use, querying such data sets using languages such as XQuery or XPath is considered truly difficult still. Therefore, without an abstraction that flattens the nested structure of XML, the natural language processing (NLP) interfaces to such databases [65, 47, 48] do not appear to have a serious influence. As argued in works such as BioBike [17, 16], for over 40 years biologists have voted overwhelmingly not to embrace computer programming as a basic tool through which to look at their world that led to BioBike's development (and many other that followed, e.g., [62, 24, 52]) that tried to use NLP as a means to gain access to needed information from the vast array of Life Sciences repositories. But the reality is that such a high level interface still remains illusive mostly because of the representation hurdles, highly interpretive nature of Life Sciences data[1], translation from natural language to database queries, and the inherent difficulty in processing NLP queries. A practical approach, we argue, is providing a flat relational view of complex data such as XML so that users can comprehend and query the information repositories at their disposal in the well understood model of flat relations as a middle ground[2]. Once chosen, the question remains, how do we facilitate comprehension of

[1]Tool applications are essential before an understanding can be gained for most biological data such as DNA sequences, protein structure, or pathways. Simple read off of the data as in relational model does not reveal any information in general.

[2]This is not to say that other data models are not useful,

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

a natural language query in a given context, device a strategy to compute its response, and implement the strategy as a traditional query over structured relations stored in local or remote information repositories?

In this paper, our goal is to show that a prudent design of multiple abstraction layers to separate the lower level structural details from the upper conceptual level aids in developing a sophisticated reasoning mechanism to analyze the semantic characterization of database objects and their inherent relationships. Once the natural language query's (NLQ) structural characterization is completed, a rule reasoner can be used to analyze the layered representation to map the NLQ to a structured query that not only implements the intent of the upper level query, it does so in a cooperative manner that no traditional query languages such as SQL or XQuery do. By that we mean, we can now actually respond to queries in much the same way a human might do and thereby reflect a deeper semantic understanding of the queries and their true intents. The diagram in figure 1 shows a bird's eye view of the interrelationships between the components in different abstraction layers, the machineries involved, and the conceptual model how the entire system functions to produce near human response to natural language queries.

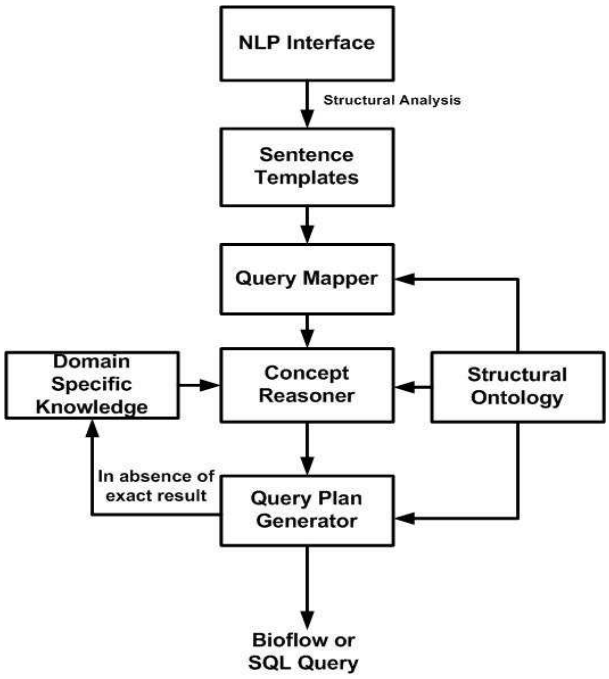

Figure 1: System components.

As a next logical step, we believe it is possible to extend this system toward a cooperative information repository that is capable of sensing the context of a query and carrying the context over an entire investigative session. To that end, we also envision a system where concepts are learned in a query session, and the interrelationship among them and the constraints they impose can be made to influence the candidate set of responses. In an early research on SmartBase [67], we have shown that non-monotonic inheritance of constraints (SQL's `where` clause conditions) related to concepts in a query to the next query as a logical conjunct simulates the possible world semantics of databases [26, 44] in a practical and computable way yet offers a better and richer approach to interactive query processing and cooperative response than systems such as [56, 49, 10] or earlier research on cooperative query processing [12, 58, 20]. We claim that contextual and knowledge rich NLP query processing remains as an extremely difficult proposition in contemporary approaches and the NLP model we propose in this paper can be adapted to include SmartBase query processing features to support a truly contextual NLQ interface.

1.1 A Motivating Example

Consider a protein/gene function database consisting of two base tables – *UniProt Protein Table* and *Entrez Gene Table* as shown in figures 2(a) and 2(b) respectively. Now consider answering queries of the form

1. *"List all F-box domain protein 2 sequences"*, or
2. *"What are the functions of UniProt proteins Q9UKT8 and Q9NVA1"*?

In a conventional relational database, these two natural language queries will most likely map to the following to two SQL queries to return the answer sets shown in figures 2(c) and 2(d) respectively.

```
select DNASequence
from Entrez as e, UniProt as u
where u.ProteinID=e.UniProtProteinID and
  ProteinName="F-box domain protein 2"

select Function
from UniProt
where ProteinID in {Q9UKT8, Q9NVA1}
```

The scientific fact here is that the UniProt protein `Q60584`, with corresponding gene name `FBXW2`, is a mouse protein. Functionally, `Q60584` is identical to UniProt protein `Q9UKT8` and have the same gene name, `FBXW2`, is a human protein. Furthermore, they also have almost identical DNA sequences. This similarity can be computed – (i) by homology of the two sequences corresponding to the UniProt IDs `Q60584` and `Q9UKT8` as shown in Entrez Gene Table, (ii) knowing that these two genes are orthologs, (iii) by determining that they have identical gene names in Entrez Gene Table, (iv) by consulting a ID mapping database such as GeneCard [1] and finding their equivalence, or (v) by consulting the Gene Ontology (GO) database [21], among many other ways. If such background knowledge is used, then it becomes possible to return the tables shown in figures 2(e) and 2(f) with an additional response in each table, as knowledge derived responses. Notice that such inferences are not directly possible in traditional relational databases, but these responses are valid, and useful. The BioBike system actually takes this approach.

In contrast, in BioBike, users must specifically perform a homolog computation in order to receive the responses shown in figures 2(e) and 2(f). The system only provides the computational means. The novelty of our system is that

but to emphasize that flat relations make it possible to use powerful deductive query engines such as Prolog, XSB [64] and F-Logic [42] without much effort as we demonstrate in this paper. In fact, the ORAKEL system [13] has been developed entirely in F-Logic.

ProteinName	ProteinID	Function
Putative replication	O85067	Plasmid maintenance
Ubiquinol-cytochrome c reductase	Q9NVA1	Cytoplasmic vesicle
F-box domain protein 2	Q60584	Substrate recognition

(a) Protein Table "UniProt"

GeneName	GeneID	UniProtProteinID	DNASequence	...
FBXW2	30050	Q60584	CTCTTTCTTTCG ...	...
repA1	1246500	O85067	ACCCTTGGAAACCC...	...
FBXW2	26190	Q9UKT8	CTCTTTCTTTCT ...	...
UQCC	55245	Q9NVA1	TTTTGGGGCCCCAAA...	...

(b) Gene Table "Entrez"

DNASequence
CTCTTTCTTTCG ...

(c) Query 1 Response

Function
Cytoplasmic vesicle

(d) Query 2 Response

DNASequence
CTCTTTCTTTCG ...
CTCTTTCTTTCT ...

(e) Enhanced Response 1

Function
Cytoplasmic vesicle
Substrate recognition

(f) Enhanced Response 2

Figure 2: Queries over UniProt and Entrez databases.

such knowledge is derived by the system on its own and no explicit requests need to be made. In a more recent work [32, 31], it was shown that an ontology based logical query language ConLog can be used to non-intrusively capture semantic knowledge and query response generated without such explicit knowledge entailed response computation. In this paper, our goal is to show that a canonical representation, as shown in figures 3(a) through 3(f), can be used to compose queries in a logic based language such as Datalog without user intervention. More importantly, we show that natural language sentences can be classified into sentence structure templates, and these templates can be used to select appropriate query predicates to fire to compute intended responses. In the next few sections, we present a detailed discission how our approach works.

1.2 Related Research

Our research in this paper is mainly influenced by the BioBike project [66, 17] at Virginia Commonwealth University. In our own lab, we are focused on developing smart declarative querying and workflow engines for distributed and heterogenous scientific databases. We have developed a declarative query language for data integration called BioFlow [38, 39] based on which we have developed two visual interfaces called VizBuilder [33, 28] and VisFlow [55, 53, 54] to support distributed workflow queries for Life Sciences data management and querying system called the LifeDB [9]. To complement LifeDB, we have also develop an ID mapping database called the MapBase [35] and phylogenetic database PhyloBase [37, 36].

The promise of knowledge rich query interfaces in Life Sciences inspired us to explore if systems could detect query intent, especially in exploratory settings such as in SmartBase in which we developed a contextual query processor capable of recognizing the query skyline [43] from a set of successive queries. Augmenting structured databases with semantic knowledge for the purpose of knowledge rich queries have been studied in the context of NLQs [32, 31], and structured database integration [34]. Despite many such efforts, articulating application description by domain scientists in Biology to facilitate knowledge rich queries remain a significant challenge. This, we believe, is mainly due to a complex regiment of interpretive tools needed to analyze the raw data, the domain specific knowledge that leads to many alternative and approximate response to a query, and so on. From these standpoints an NLP interface seems to be very attractive for this community and BioBike stands out as an interesting model as it demonstrated its ability to respond to Biological queries in an intuitive way. In this paper, our goal is to fuse these two approaches to be able to respond to natural language queries over arbitrary biological databases.

While interesting in its approach, BioBike still seems to have encoded the query processing logic into the system in a way that did not support good enough abstractions that made it a less likely candidate for a viable architecture that designers could just adopt as a generalized engine for widespread use. To a large extent, it still remained a seriously hard wired programming application. Though NLP interfaces for traditional structured databases are extensively studied in the literature (e.g., [65, 47, 48]), knowledge rich queries using NLP [40, 41] have been rarely and tangentially explored. Mostly existing NLP systems map NLQs to structured queries only when the database structures, application domain and the interpretation of the data are known and well understood. Although NLP is being increasingly used to return possible answers [14, 6, 22], we stand out in the way we conceptualize the database, represent them in a canonical form, and allow an intelligent and user specific analysis of the content of the database to map the NLQ to a target declarative query language. We are unaware of any system that follow this approach including BioBike.

The remainder of the paper is organized as follows. In section 1.1, we have already presented an overview of the system and discussed its various components on intuitive grounds. Then to illustrate the functionalities of our system, we discuss several expository examples and their execution in reference to the system components just presented in section 2. In contrast with many NLP based query processing system, we aim to support description based query processing in which a set of sentences together describe a complete workflow of an application in the direction of SmartBase. The type of NLP queries we support are introduced in section 3 and demonstrate that the type of queries we support covers a wide class of queries, and they can be nested to

ConceptA	ConceptB
Gene	Protein

(a) Derivatives: der

TableX	TableY	ColumnX	ColumnY
Entrez	UniProt	UniProtProteinID	ProteinID

(b) Foreign Keys: forK

ConceptX	ConceptY	Relation
Gene	Gene	Ortholog, Paralog, Duplication

(c) Similar Concepts: simCon

Relation	Operation
Ortholog	BLAST, ORSCAN
Paralog	GENCODE

(d) Tools: cTool

Concept	PrimaryKey	AttributeName	AttributeValue
Gene	1246500	GeneName	repA1
Gene	1246500	UniProtProteinID	O85067
Gene	1246500	DNASequence	ACCCTTGGAAACCC...
Gene	1246500	GeneID	1246500
Gene	55245	GeneName	UQCC
Gene	55245	UniProtProteinID	Q9NVA1
Gene	55245	DNASequence	TTTTGGGGCCCCAAA...
Gene	55245	GeneID	55245
Protein	O85067	ProteinID	O85067
Protein	O85067	Function	Plasmid maintenance
Protein	O85067	ProteinName	Putative replication
Protein	Q9NVA1	Protein ID	Q9NVA1
Protein	Q9NVA1	Function	Cytoplasmic vesicle
Protein	Q9NVA1	ProteinName	Ubiquinol-cytochrome c reductase
...	...	...	...

(e) Canonical Database: canDB

ConceptName	TableName	Key	Attributes
Gene	Entrez	GeneID	GeneName, DNASequence, ...
Protein	UniProt	ProteinID	Function, ProteinName, ...

(f) Structural Ontology (sO)

Figure 3: Internal representations and meta data for knowledge rich queries.

pose more semantically rich queries. Our goal is to build a middleware that is capable of interpreting the problem description by recognizing the intent of the overall workflow, and offer a near NLP experience and return a more cooperative response than traditional databases. From this standpoint, we also highlight the salient features of our system that distinguishes it from the leading contemporary systems. In section 4, we discuss the heart of our system BioSmart, its concept reasoner that helps compute knowledge rich queries using domain specific knowledge. The method to incorporate application specific desktop or online computational tools in Datalog are also discussed since custom designed and generic computational tools play a significant role in biological data processing. In section 4.4, we discuss Java specific calls to facilitate invocation of computational tools from XSB, the reasoning platform we use in BioSmart. In section 5, we discuss a more generic and more powerful approach to modeling external function calls in the form of a database engine using workflows. Finally we summarize and discuss possible future research in section 6.

2. BIOSMART SYSTEM OVERVIEW

Representation and interpretation of biological data are complex, and analyzing them requires domain specific knowledge which may vary from user to user. To support such a fluid and knowledge driven querying environment, we offer a free-from NLP interface called *BioSmart*. In this interface, users may ask arbitrary queries that are parsable as a valid natural language sentence. We then categorize these sentences into several classes that fit into predefined syntactic templates[3]. These templates are then analyzed in the context of the underlying database structures and and assigned an interpretation entailed by its logical meaning. We expect the analysis to yield a query and a query execution plan in a declarative language such as SQL, Datalog, or BioFlow.

2.1 Components and Overview of the System

Let us introduce the architecture of the NLQ engine we envision with the help of an example. Consider the query Q below.

> *"Find the photosynthetic genes of cyanobacteria Prochlorococcus sp. strain (known also as MED4)".*

Prochlorococcus is an extremely small Chl *b*-containing light-harvesting cyanobacterium antenna system sometimes constituting up to 50% of the photosynthetic biomass in the oceans [25]. So, the query above is interesting at many levels. First, although there is a fairly recent database on cyanobacteria called *CKB* [60], it is not so simple to dig out the information this query seeks from this database. The search term Prochlorococcus does not pull up any information from CKB. But, from the literature [25] we know

[3]We defer the discussion on the specific mapping algorithm from natural language to categorized templates in this paper for the sake of brevity. Instead, we mainly focus on how such templates once generated are interpreted to respond to queries.

that Prochlorococcus has two strains, `MED4` and `MIT9313` that are representatives of high and low-light adapted ecotypes, characterized by their low or high Chl *b/a* ratio, respectively. Furthermore, `MED4` is more recently evolved and has about 1,686 protein coding genes while `MIT9313` belongs to the most deeply branching lineage of Prochlorococcus with 2,200 genes. The Prochlorophyte Chlorophyll-Binding proteins (PCBs) responsible for photosynthesis in Prochlorococcus are encoded by a single gene in all the low *b/a* strains, whereas multigene families have been found in several high b/a strains. It is also known that `pcb` is a gene in the high light-adapted `MED4`, and `pcb1` and `pcb2` are two genes in low light-adapted `MIT9313` strain for photosynthesis.

However, it is not possible to isolate these genes from CKB using `pcb` as the search term to know if their function is photosynthesis. To discover this knowledge, one must sift through to select all rows having `pcb` as a column, and link the gene names to either UniProt of KEGG databases to see if their functions include photosynthesis. Two such gene names are `PMM0627`, and `pcbA` or `Pro0783` that are listed to have photosynthesis as their functions in KEGG and UniProt databases respectively. So the innocent looking query is not that simple at all to compute without knowing all these details from the literature in the first place and thus obviating the need for querying. Even when one wants to learn more, starting off with a wrong database or inappropriate search term may throw the search in a wrong direction. Finally, phylogenetically the closest relative of Prochlorococcus is the organism Synechococcus, and there is great deal of information on this organism that can be leveraged to learn details about photosynthetic genes of Prochlorococcus indirectly. Therefore, we believe the novelty of BioSmart is that given a query Q, a database D and a knowledgebase K, it computes the response to the query as the entailment relation $D \cup K \models_Q A$, which to our knowledge no other leading databases or query answering system in Life Sciences do.

2.2 The Generative Process

The entailment relation we have introduced above is captured in the schematic architecture of BioSmart shown in figure 1. The natural language sentences in BioSmart are accepted and parsed by its *NLP Interface* (NLI) and mapped to predefined sentence or query templates. It is possible that a sentence will be broken down into several such templates that capture the meaning. The *Query Mapper* (QM) component transforms the templates into a logical query using the information in the *Structural Ontology* (SO) in the context of the underlying database – identifies the tables, the analysis tools needed, and the possible joins needed to compute the query.

In the SO, we maintain concepts and their properties independent of their structural affiliation in tables or their specificity as objects in a table. For example, the two tables *Entrez* Gene Table and *UniProt* Protein Table in figures 2(a) and 2(b) will be described in SO as a meta-data table as shown in figure 3(f). The concepts in SO (e.g., Gene and Protein in column *ConceptName*) can now be linked to natural language concepts independent of their table affiliation, which now can actually be discovered from the table in figure 3(f) using column name *TableName*. Other components such as *Query Mapper* (QM), *Concept Reasoner* (CR) and *Query Plan Generator* (QPG) also use the services by and the information maintained in SO.

2.3 Query Mapping

The CR subsystem uses a set of axioms to discover how responses can be generated using the objects in the database for a given natural language sentence templates. This is the component that also discovers relationships that can be used to construct alternative responses, either directly or indirectly. It also discovers any need for tool applications to find query responses. The interesting aspect of this component is that the axioms used in CR are not domain specific, and so, the system can be used for other domains as is. However, *domain/concept specific knowledge* (DSK) in the form of a knowledge base can be supplied as a plug-in to tailor the functionality and bring specificity to the system. Better or more richer the knowledge-base is, more sophisticated response it is capable of generating. The union of DSK, CR and SO serves as the knowledgebase K in the entailment relation $D \cup K \models_Q A$, and the response A is as rich as entailed by K.

The query specification generated by the CR is then forwarded to the QPG which in consultation with the SO transforms the specification into a set of database specific executable queries in a language of choice. QPG achieves this goal by using the *canonical representation* of the existing data sources. In the canonical representation, we represent the underlying database in a triple form similar to RDF [4]. All the tables from the user-specified database are broken up in a ⟨*concept, attribute type, attribute value*⟩ triple format. Concept is a combination of the concept types and the primary key for a conceptual object. Attribute type is the column name from the original data table and attribute value is the value stored in that column for a particular concept. Canonical representation of the data sources shown in the figures 2(a) and 2(b) is presented in figure 3(e). The steps involved in the mapping process is outlined in algorithm 1 at a very high level.

Algorithm 1: Mapping NLQ Q to executable knowledge rich query Q_e.

Input: A natural language query Q
Output: Executable declarative query Q_e
Perform syntactic analysis of Q to generate a parse tree T;
Perform structural analysis of T to match a sentence template $s \in S$;
Generate logical equivalent L of s;
Apply DSK to L to generate conceptual plan C;
Generate executable script Q_e from C;
Execute Q_e;
if Q_e *succeeds* **then**
  Return result A;
  Exit;
else
  Try alternate mapping, if possible;

3. QUERY TYPES

Since interpretation of arbitrary natural language sentences are difficult, keeping in mind that most flat and structured databases use a set based model, natural language sentences that mimic `select-project-join` (SPJ) queries are our priority. Such queries have limited interpretive scopes in terms of what they allow. For example, queries are mostly about objects and their properties, and their relationships

with other objects. Often, we need to construct complex objects by piecing together parts from various tables. We believe almost all such query structures can be admitted if we recognized three basic types of sentence structures – *iterative*, *conditional*, and *imperative* or *interrogative* natural language queries. Keeping the admissible sentence structures set simple and small, we actually follow SQL's footsteps in which complex and more expressive queries can be built by nesting simple structures arbitrarily. Thus, most of the biological queries potentially can be expressed by one of these types, or by a combination of them. In BioSmart, we do so by allowing series of queries in succession in a context.

3.1 Iterative Queries

Consider expressing a query in English in two different but semantically equivalent ways.

"List the functions of all human genes"

and,

"For all Homo sapiens genes, list their functions."

These two seemingly two structurally different queries in English actually map identical structured queries in SQL, or at least can be expressed by a single query from table 2(a). But as a query type, these natural language queries ask to retrieve objects that satisfy certain properties (including empty properties). If we parse these queries in a syntax tree as shown in figure 4, we will discover that each of these query types roughly adhere to one of the parse tree structures in this figure, we call sentence templates. The queries above can be parsed to resemble the template shown in figure 4(a), or the iterative type.

An iterative, or *loop*, type natural language query, and thus its template, starts with a prepositional phrase (PP), followed by a noun phrase (NP) and then a verb phrase (VP). The prepositional phrase usually includes the phrase "for all", or its variants. The noun phrase essentially points to the objects or entities we are to apply the verb phrase, i.e., properties or actions. The parse tree for the query

"For all genes of cyanobacteria find homologs"

is shown in figure 5(a) which actually is an instance of the iterative query template in figure 4(a). In this query the NP is *all genes of cyanobacteria* and the VP is *find homologs*.

3.2 Conditional Queries

Iterative queries are not required to satisfy any constraint. In other words, they can be viewed as a simple projection query in SQL (i.e., `select from`), possibly with a simple `where` clause condition for the purpose of object or property identification. In contrast, a *conditional* query specifies an arbitrary precondition that the objects must satisfy in the form of an *if then* structure as shown in the query template in figure 4(b). In such queries, a NP-VP sequence follows the *if* condition (i.e., the S' subtree), and the VP following the NP captures the action clause, where the NP-VP sequence has a simple *subject-verb-object* sentence structure. The VP after the *then* has a *verb-object* imperative sentence structure. Figure 5(b) shows the parse tree for the query

"If gene UQCC is protein coding, then find its protein",

as an instance of the template in figure 4(b). In this example *"is protein coding"* is the VP for the *if* condition, and *"find its protein"* is the VP for the *then* inference.

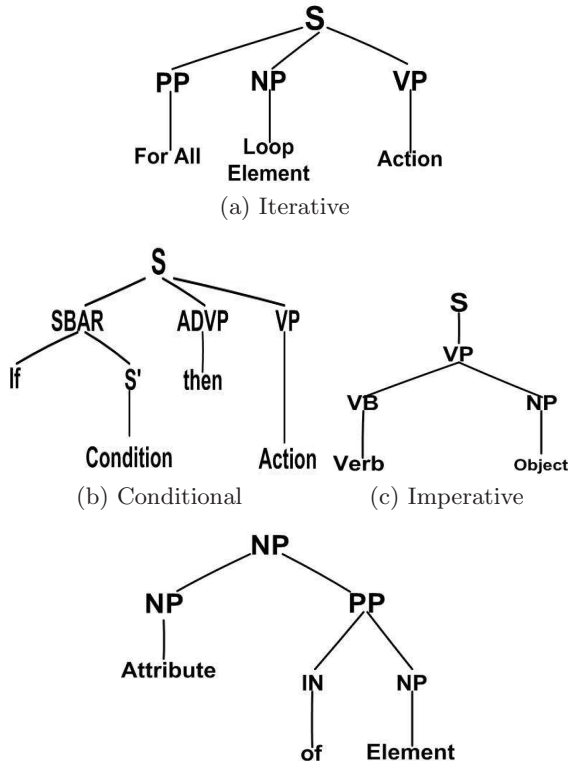

(a) Iterative

(b) Conditional

(c) Imperative

(d) The noun phrase structure

Figure 4: Natural language query templates.

3.3 Imperative Queries

Imperative sentences or queries are basically a verb phrase (VP) consisting of a verb (VB) and an object (NP) on which the verb is to be applied. The template representing an imperative query is shown in figure 4(c). These type of sentences are used as standalone queries or as a part of more complex iterative of conditional queries. Usually the object (NP) has a structure of the form *"attribute of element"* as in figure 4(d). As an example, consider the query

"List all genes of cyanobacteria",

and its parse tree shown in figure 5(c). Here the verb is *List* and the object is *genes of cyanobacteria* in conformance with the structure of figure 4(d). As discussed in section 2.1, we will be correct to return genes of Prochlorococcus and Synechococcus, and `MED4` and `MIT9313`, such as `pcb` and `pcbA`.

4. KNOWLEDGE RICH QUERYING

The query identification and processing apparatus we have introduced earlier can now be leveraged to respond to the simple knowledge rich query

"Find the function of gene repA1"

in Escherichia coli. Though this query appears simple, as discussed in section 2.1, finding answer to it may require the use of multiple biological data sources and tools. A user can go to NCBI GenBank to find its function. Failing to find the answer in GenBank, she can find its UniProt ID `P03066`

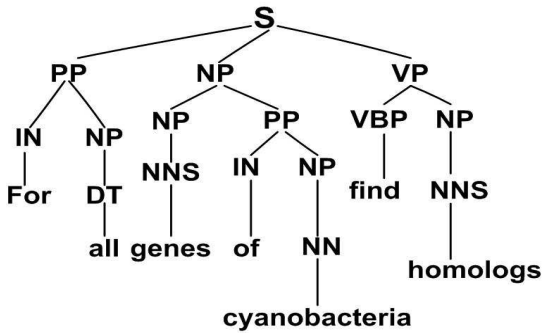

(a) Iterative

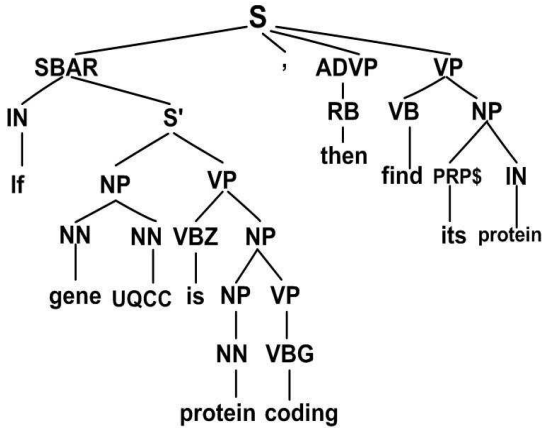

(b) Conditional

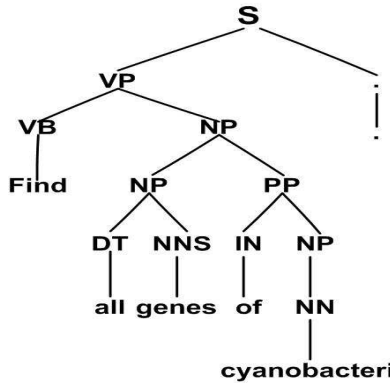

(c) Imperative

Figure 5: Natural language query examples.

and try to find its function from UniProt database, and a full annotation can also be obtained from GO database. In UniProt, the function for `repA` (Replication initiation protein) is listed as `plasmid replication`, and `copy control`, whereas GO annotation includes {`DNA replication, plasmid maintenance`}. In the event none of these databases produced any useful information, she could try to use NCBI BLAST search to find the orthologs of `repA1` and then find the function of those orthologs.

Such approaches require a user to be familiar with all these details, and complex queries require even more complex and intricate interwoven knowledge will be essential. Moreover, different resources and tools have unique representation and naming policy which makes it even more difficult to find answers. For example, the GeneCards database assigns `RPA1` as the ID for `repA1`, which is supposed to be an authoritative site for finding the IDs for `repA1` in various databases. Apparently, `repA1` gene symbol has been discontinued and replaced with the name `RPA1`. To aid biologists in such a confusing landscape, in BioSmart, we approach the computation of the query in several steps. We use the Stanford parser to first parse the query submitted to BioSmart NLI and attempt a mapping to one of the query templates in figure 4 as resulting in the imperative query parse tree in figure 6, with action *Find* on attribute *function* of element *gene repA1*.

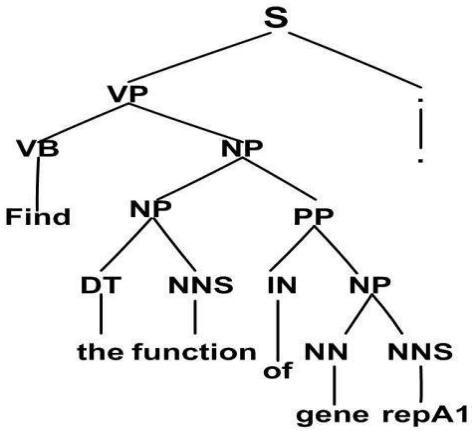

Figure 6: Parsed tree of example query.

4.1 Direct Concept Reconstruction

From the discussion in section 1.1 it may not have been apparent that we actually do not use the base tables in figure 2. Instead, we always use the derived tables in figure 3. In particular, all tables are collapsed into the canonical database form in figure 3(e) that allows reconstruction of the objects in any table using an object identifier, the primary key of each table represented as a concept. The remaining set of tables in figure 3 forms the DSK and SO components. The reconstruction process depends on the query type, and the complexity of the query and its interpretation assigned.

In the query above, the goal is to compute the *function* (a property) of a *gene* (a concept/object) named *repA1* (attribute value) corresponding to another property (i.e., *gene name*). QM uses this interpretation to map the properties (to the attribute *AttributeName*, the concept to the attribute *Concept*, and attribute value to *AttributeValue*. Essentially, we are trying to compute the pair ⟨*function, AttributeValue*⟩ for the concept *gene* having ⟨*GeneName, repA1*⟩ as an entry in the canonical database *CanDB*. Logically, the response can be constructed by an equi-join of the *CanDB* table on the column *PrimaryKey* for the *Concept gene*. In general, the reconstruction is captured in the following CR rule,

```
1: res(Con, Pk, AttName, AttVal) :-
      canDB(Con, Pk, AttName, AttVal).
```

and with the conjunctive query below that is equivalent to the equi-join above.

```
? res('Gene', Pk, 'GeneName', 'repA1'),
    res('Gene', Pk, 'Function', Val).
```

If k number of attributes of a concept were needed to be computed, we would be required to write a k-way conjunctive query in this approach.

Obviously, this query will fail. But, had the query been,

```
? res('Gene', Pk, 'GeneName', 'repA1'),
    res('Gene', Pk, 'UniProtProteinID', Val).
```

it would have succeeded producing a binding `O85067` for *Val*. It failed because the concept *gene* does not have a *function* in the base table *Entrez*, and consequently in the table *CanDB*. Therefore, the response remains empty unless CR tries an alternative evaluation.

4.2 Indirect Response Generation

The query above failed because concept *gene* does not have a property called *function*. Biologically, we know that genes encode proteins, and thus gene functions are manifest as protein functions, and thus are synonymous. Such knowledge are captured in the table *der* in figure 3(a). A transitive closure of this relationship captures what we can compute as substitutes for a given concept, as captured in the axioms in rules 3 and 4 below.

```
2: res(Con, Pk, AttName, AttVal) :-
      rel(Con, Der, Pk, PkD),
      res(Der, PkD, AttName, AttVal).

3: rel(Con, Der, PkC, PkD) :- der(Con, Der),
      sO(Con, TabC, KeyC, AttsC),
      member(ColC, AttsC),
      sO(Der, TabD, KeyD, AttsD),
      member(ColD, AttsD),
      forK(TabC, TabD, ColC, ColD),
      canDB(Con, PkC, ColC, ValC),
      canDB(Der, PkD, ColD, ValC).
4: rel(Con, Der, PkC, PkD) :-
      rel(Con, DerC, PkC, PkDC),
      sO(DerC, TabC, KeyC, AttsC),
      member(ColC, AttsC),
      sO(Der, TabD, KeyD, AttsD),
      member(ColD, AttsD),
      forK(TabC, TabD, ColC, ColD),
      canDB(DerC, PkDC, ColC, ValC),
      canDB(Der, PkD, ColD, ValC).
5: member(Mem, [Mem|_]).
6: member(Mem, [_|Tail]) :- member(Mem, Tail).
```

To reconstruct the extended object, for the property pair ⟨*function, AttributeValue*⟩ from the corresponding concept `Protein`, we need to establish the fact that `Protein` indeed is a substitute (in the *der* table in figure 3(a)), *UniProt-ProteinID* `O85067` corresponds to *GeneID* `1246500` (in table *Entrez* in figure 2(b)), and that the *Protein* `O85067` has the pair ⟨*Function, AttributeValue*⟩ in table *CanDB*, by virtue of table *UniProt* in figure 2(a). We do so in rule 2 above by asserting a transitive relationship between the concepts `Gene` and `Protein`, and making sure they are connected by a foreign key relationship.

The rules 3 and 4 are a bit involved, but are conceptually simpler. These rules basically say, two objects are connected by a derived relationship is they have a direct foreign key relationship (rule 3), or a transitive foreign key relationship between the concept C and S such that C has a direct foreign key relationship with some concept D in *der* table and a foreign key relationship between the derived concept D and some other concept S. The rules 5 and 6 are necessary axioms to test list memberships used in rules 3 and 4 to inspect the base tables have the required attributes. Adding rules 2 through 6 will now help us evaluate the subquery

```
res('Gene', Pk, 'Function', Val)
```

to be true with a binding `Plasmid maintenance` for the variable *Val*, essentially computing the *Function* for the `gene repA1`, via *ProteinID* `O85067`, and *GeneID* `1246500`.

4.3 Interpretive Queries

The indirect responses in the previous section are actually a class of queries that use conceptual substitutions to derive responses, i.e., property of proteins for genes know that they are substitutable as captured in the table *der* as a derivative. In fact any such biological knowledge can be encoded in BioSmart by creating a new rule for the predicate *rel/4* to link objects or concepts through their identifying keys. The independence of the *rel/4* predicate from how the antecedent is represented, makes it possible to make the mapping algorithm abstract and generic. For example, Scandinavian males, especially 33%-45% Swedish males, carry the I1 haplotype. We could then link genes of an offspring to his geographic origin, or to someone in that region, to discover possible properties. In such cases too, the predicate *rel/4* can be leveraged to link seemingly unrelated pieces of information to derive knowledge.

But such relationships are required to be in one of the database tables as ground values. In other words, no new information is actually generated, they are only linked in a meaningful and informative way. In BioSmart, we do allow a third kind of queries that actually allows generation of new knowledge not available in the database as ground facts using computational tools or functions. Applications of such tools have the potential to reveal new relationship among the database objects previously unknown, or discover new properties of objects in it. As shown in figure 7, species, as well their genes and morphologies, are related via genetic and morphological homology such as orthologs, paralogs, and horizontal gene transfer.

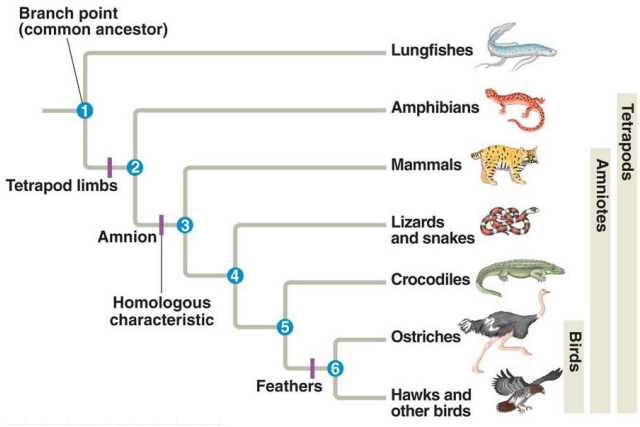

Figure 7: Example phylogenetic tree.

Aside from storing base facts in tables such as in PhylomeDB [29] or TreeBASE [71], we can actually apply computational tools to compute these relationships and infer

new properties. For example, orthologs can be computed using BLAST like tools such as Ortholog-Finder [27], ORCAN [73] or GENCODE [59], and phylogeny construction tools such as MEGA [45], PAUP [68] or PHYLIP [50]. New relationships can also be established using databases such as GeneCards [63], and MapBase using ID mapping. In BioSmart, we allow such tool application based on background knowledge to link objects and then use the direct and indirect construction of properties using discovered relationships. For example, for the gene *repA1*, there are several orthologs that can be computed or retrieved from Entrez database or by using tools such as Ortholog-Finder or ORCAN: Gene ID `327491` in zebra fish, Gene ID `68275` in mouse, and Gene ID `417563` in chicken, each one of which could be used to decipher properties for the gene `repA1`. It is interesting to note that recent studies show that although `repA1` does not have paralogs in animals, it has paralogs in plants: `RPA1A` to `RPA1E` in Arabidopsis thaliana, and `RPA1B` in Glycin max, `RPA1C` in Sorghum bicolor, and so on [5]. This result is only available as of today in scientific literature, and not in any database, indicating that a text mining tool such as GeneView [69] may be appropriate to discover this knowledge.

Calling such external functions from logic based languages such as Datalog, Prolog or XSB, is application specific. In general, they are called *foreign* codes or functions, and have specific protocols for implementation. In BioSmart, we use XSB [61] as a reasoning platform for its set based processing strategy, and Interprolog [11, 2] Java API for procedure calls. Interface *javaMessage()* of Interprolog binds XSB predicates with Java procedures. The rule implementing interpretive queries thus take the form below.

```
7: res(Con, PkC, AttNameR, AttValR) :-
      res(Con, PkC, _, _), simCon(Con, RelCon, Reln),
      cTool(RelnO, Ops), member(RelnO, Reln),
      res(RelCon, PkR, AttNameR, AttValR),
      member(Op, Ops), applyOp(Op, PkC, PkR).
```

The rule 7 above basically asserts that we can derive the pair $\langle AttNameR,\ AttValR\rangle$ for a object *Con* with a identifier *PkC*, if it is related to another object via the relationship in *simCon* table and an external call to the tool *Op* can verify the stated relationship with object *RelCon* with identifier *PkR* having the pair $\langle AttNameR,\ AttValR\rangle$ as its property.

4.4 Computational Tool Integration with XSB

Predicate **applyOp()** in rule 7 above is a foreign function for XSB which is implemented procedurally in Java. It invokes different types of functions depending on the argument *Op*, i.e., it may access databases such as PhylomeDB, TreeBase, GeneCards or MapBase, it may initiate a computation by running desktop tools such as MEGA, PHYLIP, or PAUP, or look up the information in an web accessible tool such as WebPHYLIP or ORCAN. We use the *Interprolog* [2] API to switch between the XSB reasoner and the procedural Java environment. Interface *javaMessage()* of Interprolog binds XSB predicates with Java procedures. The general form of *javaMessage()* interface is shown bellow.

```
javaMessage(Target, Result, Exception, MessageName,
   ArgList, NewArgList)
```

Interface *javaMessage()* synchronously calls a method of a Java object *Target*, then waits for its *Result*, catching any *Exception* that may occur. *ArgList* specifies the arguments necessary for the function call, which must be of the proper Java-compatible types. *NewArgList* contains the same objects in *ArgList* reflecting possible state changes after the function has been processed. Here, *MessageName* is the invoked method of object *Target*.

A Java class called *MessageCaller* has been implemented to model the functions of *javaMessages* invoked by XSB. A method named *performOperation()* of *MessageCaller* calls online tools such as BLAST, ORCAN, WebPHYLIP, or GENCODE based on the operations specified (sent as the argument list *ArgList* by *javaMessage*). For example, to compute homology of genes by the online tool BLAST, method *performOperation()* uses the *NCBI BLAST Java Interface* [3]. Given the ID of a gene, BLAST returns the IDs of the orthologs as an XML file. We then parse the XML file to retrieve the *GeneIds* of the orthologous genes and send them back to XSB as *Result*. An alternative technique for sending results is to store them in a database table and use an XSB predicate to retrieve those results within the XSB reasoner.

To be precise, the **applyOp**(*Op, PkC, PkR*) call in XSB is handled by the *javaMessage()* operation as follows. XSB initiates the *javaMessage()* call with the instantiations *Target* to *MessageCaller*, *Result* to *PkR*, *MassageName* to *performOperation()*, and *ArgList* to $\langle Op,\ PkC\rangle$. *Exception* and *NewArgList* are returned by Java as appropriate. Note that, the argument *Op* depends on the adornment of the variable by XSB before the call and depends on the tool list *Ops* in the *cTool* predicate. Eventually, method *performOperation()* of class *MessageCaller* applies the appropriate tool to perform the intended operation expected in the **applyOp**(*Op, PkC, PkR*) predicate.

5. ACCESSING ONLINE TOOLS AND DATABASES

In BioSmart, we leverage another level of abstraction for the interpretive class queries that opens up the possibility endless ways computing them, and infer knowledge in unprecedented ways. That also means the Java *MessageCaller* will need to be implemented on a case by case, which we believe is a daunting task. To avoid such unique implementation for each call type, we have decided to use the power of the abstraction encoded into the BioFlow language introduced earlier. This language, and its variants, have been leveraged in our implementation of LifeDB, MapBase, VizBuilder and VizFlow. In BioFlow, tools, databases and online web interfaces are viewed as function calls and are abstracted uniformly. Therefore, depending on the invocation and the specific tool, it is capable of customizing the evaluation.

In most tool applications, database processing and web applications, some form of selection conditions are applied (input arguments to a function), and results are extracted (output of the operation) before and after the operations are performed. There is also some form of schema mismatch between the terminologies used in the XSB program and the target system. Without the abstraction, users and the Java application writers will need to be fully aware of these terminologies and resolve the disparities manually. The BioFlow syntax already includes the machineries needed for schema matching and wrapping, in addition to accessing remote sites. The statement to access deep web resources in

BioFlow is the extract statement with the following syntax,

extract $A_1, \ldots, A_k$
using matcher μ wrapper ω filler ϕ
from φ submit r where θ

where θ is the form condition, $A_1, \ldots, A_k$ is the projection list, μ is the schema matcher (e.g., PruSM [57]), ω is the wrapper (e.g., FastWrap [7]), ϕ is the form filler (e.g., iForm [70]), and φ is the web form address or the form function. Note that this statement returns a table by submitting columns from each row in r to the deep web database at φ. When a matcher, wrapper or filler isn't necessary, the corresponding clauses can be omitted. This statement can be constructed using the stepwise mapping of a resource such as d_1. Note that, in the case of a web service at φ, a form filler and a wrapper are not needed as the web service itself handles these functions, and no form conditions (θ) are necessary either. But a transformation function τ may be required to convert XML to flat list of attributes, which could be made available in BioSmart library and help BioFlow extract the target fields automatically with a statement similar to the one below.

extract $A_1, \ldots, A_k$
using matcher μ transformer τ
from φ submit r

Analogously, when the text data are delimited in some way and are flat, we can expect a statement similar to the one below when inputs are not needed and hence the submit clause can be omitted, but a wrapper ω is required to be able to read the data.

extract $A_1, \ldots, A_k$
using matcher μ wrapper ω
from φ

These features of BioFlow allow users to write applications without having to worry about the details of extraction methods, location, technology specific nuances, format, and schema heterogeneity. The adoption of BioFlow also allows us to implement complex workflows using the applications similar to VisFlow. Users now can conceptualize a complete workflow at the highest abstraction level keeping a global scheme in mind knowing that the underlying data management and integration apparatus will be able to map her application onto potentially heterogeneous resources correctly and efficiently without any loss in query semantics.

For the call **applyOp**(*BLAST, repA1, PkR*), we will construct the following BioFlow statement to execute.

extract *GeneID*
using matcher *PruSM* wrapper *FastWrap* filler *iForm*
from *'http://blast.ncbi.nlm.nih.gov/Blast.cgi'*
submit *ArgRel*

In the above expression, we will supply `repA1` as the lone tuple in the table *ArgRel*. In fact, BioFlow can process set of inputs and thus is capable of a set based processing which can be utilized to speed up XSB predicate evaluation using external functions.

6. SUMMARY AND FUTURE RESEARCH

While there is a great deal of opportunities and interests in non-traditional data management and querying using key word based, NoSQL or natural language over unstructured databases, querying of structured databases using these approaches are also equally interesting. In this paper, and in our earlier research, we have demonstrated that interesting knowledge rich queries can be answered using such approaches, especially in investigative applications. Our contention is that users need not go to unimaginable lengths to dig out information just because they did not know how to. The BioSmart system we propose demonstrates the opportunities that exist and what is possible.

The concept reasoner, and domain specific knowledgebase we have leveraged is one of the major tool boxes that make BioSmart actually smart. But designing these components are manual, often application specific and tedious. But it also makes it possible to apply BioSmart to other scientific domains just by changing these components. Opportunities exists to use ontologies such as GO and SnoMed CT [30, 72] to help users frame effective natural language queries by allowing terminology independence. This will require mapping mapping query terms to conceptual terms in the ontologies and use the standardized terms in the mapped queries. We are investigating an approach designing the CR and DSK components, at least partially, automatically from these components.

In Schema-Free SQL [48], meta information such as relation trees were used to aid query processors to tolerate mismatch or errors in schema information, and in BioVis [34] and ConLog [31] conceptual structures were leveraged to map queries properly to underlying scheme. In BioSmart the need for such a structure to appropriately select parts of schema to frame the queries is much greater. Currently, the engineering of this structure is application specific and manual. Developing a similar conceptual structure generation scheme at least semi-autonomously will significantly improve usability of BioSmart. These are some of the issues we plan to continue as our future research.

7. REFERENCES

[1] *GeneCards: The Human Gene Compendium.* http://www.genecards.org/.

[2] *InterProlog 2.1.2: a Java front-end and enhancement for Prolog.* http://www.declarativa.com/interprolog/.

[3] *NCBI BLAST Java Interface (Concordia University).* http://users.encs.concordia.ca/f_kohant/ncbiblast/.

[4] *Web Services Description Language (WSDL) Version 2.0.* http://www.w3.org/TR/wsdl20/.

[5] B. Aklilu and K. Culligan. Molecular evolution and functional diversification of replication protein a1 in plants. *Frontiers in Plant Science*, 7(33), January 2016.

[6] N. Aletras, D. Tsarapatsanis, D. Preotiuc-Pietro, and V. Lampos. Predicting judicial decisions of the european court of human rights: a natural language processing perspective. *PeerJ Computer Science*, 2:e93, 2016.

[7] M. S. Amin and H. M. Jamil. An efficient web-based wrapper and annotator for tabular data. *IJSEKE*, 20(2):215–231, 2010.

[8] Y. Amsterdamer, A. Kukliansky, and T. Milo. A natural language interface for querying general and individual knowledge. *PVLDB*, 8(12):1430–1441, 2015.

[9] A. Bhattacharjee, A. Islam, M. S. Amin, S. Hossain,

S. Hosain, H. M. Jamil, and L. Lipovich. On-the-fly integration and ad hoc querying of life sciences databases using LifeDB. In *DEXA*, pages 561–575, 2009.

[10] J. Boulos, N. N. Dalvi, B. Mandhani, S. Mathur, C. Ré, and D. Suciu. Mystiq: a system for finding more answers by using probabilities. In *SIGMOD*, pages 891–893, 2005.

[11] M. Calejo. InterProlog: Towards a declarative embedding of logic programming in Java. In *IEEE International Conference on Robotics and Automation*, pages 714–717. Springer, 2004.

[12] W. W. Chu. Cobase: A cooperative query answering facility for database systems. In *DEXA, Prague, Czech Republic, September 6-8*, pages 134–145, 1993.

[13] P. Cimiano, P. Haase, J. Heizmann, M. Mantel, and R. Studer. Towards portable natural language interfaces to knowledge bases - the case of the ORAKEL system. *DKE*, 65(2):325–354, 2008.

[14] B. L. Cook, A. M. Progovac, P. Chen, B. Mullin, S. Hou, and E. Baca-Garcia. Novel use of natural language processing (NLP) to predict suicidal ideation and psychiatric symptoms in a text-based mental health intervention in madrid. *Comp. Math. Methods in Medicine*, 2016:8708434:1–8708434:8, 2016.

[15] M. Dubey, S. Dasgupta, A. Sharma, K. Höffner, and J. Lehmann. AskNow: A framework for natural language query formalization in SPARQL. In *ESWC, Heraklion, Crete, Greece, May 29 - June 2*, pages 300–316, 2016.

[16] J. Elhai. Humans, computers, and the route to biological insights: Regaining our capacity for surprise. *Journal of Computational Biology*, 18(7):867–878, 2011.

[17] J. Elhai, A. Taton, J. Massar, J. K. Myers, M. Travers, J. Casey, M. Slupesky, and J. Shrager. BioBIKE: A Web-based, programmable, integrated biological knowledge base. *Nucl. Acids Res.*, 37:W28–32, 2009.

[18] S. Ferré. Sparklis: An expressive query builder for SPARQL endpoints with guidance in natural language. *Semantic Web*, 8(3):405–418, 2017.

[19] P. A. Fontelo, F. Liu, and M. J. Ackerman. askmedline: a free-text, natural language query tool for medline/pubmed. *BMC Med. Inf. & Decision Making*, 5:5, 2005.

[20] T. Gaasterland and C. Sensen. Using multiple tools for automated genome interpretation in an integrated environment. *Trends in Genetics*, Feb. 1996.

[21] Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nat. Genet.*, 25:25–29, 2000.

[22] D. Gkatzia, O. Lemon, and V. Rieser. Natural language generation enhances human decision-making with uncertain information. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, Berlin, Germany, Volume 2: Short Papers*, 2016.

[23] E. J. Goldsmith, S. Mendiratta, R. Akella, and K. Dahlgren. Natural language query in the biochemistry and molecular biology domains based on cognition search. *Nature*, 2008.

[24] T. Hamon, N. Grabar, and F. Mougin. Querying biomedical linked data with natural language questions. *Semantic Web*, 8(4):581–599, 2017.

[25] W. R. Hess, G. Rocap, C. S. Ting, F. Larimer, S. Stilwagen, J. Lamerdin, and S. W. Chisholm. The photosynthetic apparatus of Prochlorococcus: Insights through comparative genomics. *Photosynth Res*, 70(1):53–71, 2001.

[26] J. Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.

[27] T. Horiike, R. Minai, D. Miyata, Y. Nakamura, and Y. Tateno. Ortholog-Finder: A tool for constructing an ortholog data set. *Genome Biology and Evolution*, 8(2):446, 2016.

[28] S. Hossain and H. M. Jamil. A visual interface for on-the-fly biological database integration and workflow design using VizBuilder. In *DILS*, pages 157–172, July 2009.

[29] J. Huerta-Cepas, S. Capella-Gutiérrez, L. P. Pryszcz, M. Marcet-Houben, and T. Gabaldón. PhylomeDB v4: zooming into the plurality of evolutionary histories of a genome. *Nucleic Acids Research*, 42(Database-Issue):897–902, 2014.

[30] IHTSDO. *SNOMED-CT*. http://www.ihtsdo.org/snomed-ct. Accessed: May 14, 2016.

[31] H. Jamil. A natural language interface plug-in for cooperative query answering in biological databases. *BMC Genomics*, 13(Suppl 3):S4, 2012.

[32] H. M. Jamil. Toward a cooperative natural language query interface for biological databases. In *IEEE BIBM*, Atlanta, GA, November 2011.

[33] H. M. Jamil. Designing integrated computational biology pipelines visually. *TCBB*, 10(3):605–618, May/June 2013.

[34] H. M. Jamil. Mapping abstract queries to big data web resources for on-the-fly data integration and information retrieval. In *ICDE Workshops, Chicago, IL, USA, March 31 - April 4*, pages 62–67, 2014.

[35] H. M. Jamil. Improving integration effectiveness of ID mapping based biological record linkage. *IEEE/ACM TCBB*, 12(2):473–486, 2015.

[36] H. M. Jamil. Pruning forests to find the trees. In *SSDBM, Budapest, Hungary, July 18-20*, pages 18:1–18:12, 2016.

[37] H. M. Jamil. A visual interface for querying heterogeneous phylogenetic databases. *IEEE/ACM TCBB*, 14(1):131–144, 2017.

[38] H. M. Jamil and A. Islam. The power of declarative languages: A comparative exposition of scientific workflow design using BioFlow and Taverna. In *IEEE SWF*, pages 322–329, July 2009.

[39] H. M. Jamil, A. Islam, and S. Hossain. A declarative language and toolkit for scientific workflow implementation and execution. *IJBPIM*, 5(1):3–17, 2010.

[40] S. W. Joseph and R. Aleliunas. A knowledge-based subsystem for a natural language interface to a database that predicts and explains query failures. In *ICDE, April 8-12, Kobe, Japan*, pages 80–87, 1991.

[41] E. Kaufmann and A. Bernstein. Evaluating the usability of natural language query languages and interfaces to semantic web knowledge bases. *J. Web*

Sem., 8(4):377–393, 2010.

[42] M. Kifer, G. Lausen, and J. Wu. Logical foundations of object-oriented and frame-based languages. *Journal of ACM*, 42(4):741–843, 1995.

[43] D. Kossmann, F. Ramsak, and S. Rost. Shooting stars in the sky: An online algorithm for skyline queries. In *VLDB, August 20-23, Hong Kong, China*, pages 275–286, 2002.

[44] S. Kripke. *Semantical analysis of modal logic*. Logik und Grundlagen der Mathematik, 1963.

[45] S. Kumar, M. Nei, J. Dudley, and K. Tamura. MEGA: a biologist-centric software for evolutionary analysis of DNA and protein sequences. *Briefings in Bioinformatics*, 9:299–306, 2008.

[46] C. Lawrence and S. Riezler. NLmaps: A natural language interface to query openstreetmap. In *COLING, December 11-16, Osaka, Japan*, pages 6–10, 2016.

[47] F. Li and H. V. Jagadish. Understanding natural language queries over relational databases. *SIGMOD Record*, 45(1):6–13, 2016.

[48] F. Li, T. Pan, and H. V. Jagadish. Schema-free SQL. In *SIGMOD, Snowbird, UT, USA, June 22-27*, pages 1051–1062, 2014.

[49] Y. Li, C. Yu, and H. V. Jagadish. Enabling schema-free xquery with meaningful query focus. *VLDB J.*, 17(3):355–377, 2008.

[50] A. Lim and L. Zhang. WebPHYLIP: a web interface to PHYLIP. *Bioinformatics*, 15(12):1068, 1999.

[51] C. D. Maio, G. Fenza, V. Loia, and M. Parente. Natural language query processing framework for biomedical literature. In *IFSA-EUSFLAT-15, Gijón, Spain., June 30.*, 2015.

[52] B. McKnight and I. B. Arpinar. Linking and querying genomic datasets using natural language. In *BIBM, Philadelphia, PA, USA, October 4-7, 2012*, pages 1–4.

[53] X. Mou, H. M. Jamil, and X. Ma. Visflow: A visual database integration and workflow querying system. In *ICDE, San Diego, CA, USA, April 19-22*, 2017. To appear.

[54] X. Mou, H. M. Jamil, and R. Rinker. Visual orchestration and autonomous execution of distributed and heterogeneous computational biology pipelines. In *IEEE BIBM Shenzhen, China, December 15-18*, pages 752–757, 2016.

[55] X. Mou, H. M. Jamil, and R. Rinker. Implementing computational biology pipelines using visflow. *IJDMB*, 2017. Accepted for publication.

[56] A. Nandi and H. V. Jagadish. Qunits: queried units in database search. In *CIDR*, 2009.

[57] T. H. Nguyen, H. Nguyen, and J. Freire. PruSM: a prudent schema matching approach for web forms. In *CIKM*, pages 1385–1388, 2010.

[58] L. N. P. Godfrey, J. Minker. An architecture for a cooperative database system. In *ADBy*, pages 3–24, 1994.

[59] B. Pei, C. Sisu, A. Frankish, C. Howald, L. Habegger, X. J. Mu, R. Harte, S. Balasubramanian, A. Tanzer, M. Diekhans, A. Reymond, T. J. Hubbard, J. Harrow, and M. B. Gerstein. The GENCODE pseudogene resource. *Genome Biology*, 13(9):R51, 2012.

[60] A. P. Peter, K. Lakshmanan, S. Mohandass, S. Varadharaj, S. Thilagar, K. A. Abdul Kareem, P. Dharmar, S. Gopalakrishnan, and U. Lakshmanan. Cyanobacterial knowledgebase (CKB), a compendium of cyanobacterial genomes and proteomes. *PLOS ONE*, 10(8):1–12, 08 2015.

[61] P. Rao, K. Sagonas, T. Swift, D. Warren, and J. Freire. XSB: A System for Efficiently Computing Well-Founded Semantics. In *LPNMR*, pages 430–440. Springer, 1997.

[62] L. Safari and J. D. Patrick. Restricted natural language based querying of clinical databases. *Journal of Biomedical Informatics*, 52:338–353, 2014.

[63] M. Safran, I. Dalah, J. Alexander, N. Rosen, T. Iny Stein, M. Shmoish, N. Nativ, I. Bahir, T. Doniger, H. Krug, A. Sirota-Madi, T. Olender, Y. Golan, G. Stelzer, A. Harel, and D. Lancet. GeneCards Version 3: the human gene integrator. *Database*, 2010(0):baq020–, 2010.

[64] K. Sagonas, T. Swift, and D. S. Warren. XSB as an efficient deductive database engine. *SIGMOD Rec.*, 23(2):442–453, 1994.

[65] D. Saha, A. Floratou, K. Sankaranarayanan, U. F. Minhas, A. R. Mittal, and F. Özcan. ATHENA: an ontology-driven system for natural language querying over relational data stores. *PVLDB*, 9(12):1209–1220, 2016.

[66] J. Shrager. The evolution of BioBike: Community adaptation of a biocomputing platform. *Studies in History and Philosophy of Science*, 38:642Ű656, 2007.

[67] K. Z. Sultana, A. Bhattacharjee, M. S. Amin, and H. M. Jamil. A model for contextual cooperative query answering in e-commerce applications. In *FQAS*, Roskilde, Denmark, October 2009.

[68] D. L. Swofford. *PAUP*: Phylogenetic Analysis Using Parsimony (*and Other Methods). Version 4.* Sinauer Associates, Sunderland, Massachusetts, 2003.

[69] P. Thomas, J. Starlinger, A. Vowinkel, S. Arzt, and U. Leser. GeneView: a comprehensive semantic search engine for PubMed. *Nucleic Acids Research*, 40(W1):W585–W591, 2012.

[70] G. A. Toda, E. Cortez, A. S. da Silva, and E. S. de Moura. A probabilistic approach for automatically filling form-based web interfaces. *PVLDB*, 4(3):151–160, 2010.

[71] R. A. Vos, H. Lapp, W. H. Piel, and V. Tannen. TreeBASE2: Rise of the Machines. *Nature Precedings*, (713), 2010.

[72] H. Wasserman and J. Wang. An applied evaluation of SNOMED CT as a clinical vocabulary for the computerized diagnosis and problem list. In *AMIA, Washington, DC, USA, November 8-12*, 2003.

[73] A. Zielezinski, M. Dziubek, J. Sliski, and W. M. Karlowski. ORCAN - a web-based meta-server for real-time detection and functional annotation of orthologs. *Bioinformatics*, ePub, 2017. doi: 10.1093/bioinformatics/btw825.